

Claims 1-37 remain pending in the application, with claims 1, 18 and 31-37 being the independent claims, and with claims 35-37 having been withdrawn from consideration. Reconsideration and further examination are respectfully requested.

Claims 1-34 have been rejected under 35 USC § 103(a) over U.S. Patent 6,496,477 (Perkins) in view of U.S. Patent 6,501,756 (Katsube). Reconsideration of this rejection is respectfully requested for the following reasons.

The present invention pertains to techniques for probing a network based upon a data packet received over the network. One example of an application of the present invention is as follows. A Web server receives a data packet from a computer on the Internet. The received data packet requests that a Web page be downloaded to the requesting computer. In the data packet is a source identifier, to which the Web page will be transmitted, together with an indication of the number of "hops" that the packet made when traveling from the source computer to the Web server (typically, the TTL field of the data packet).

Rather than simply sending a generic Web page in response to this request, the present inventor has discovered that it often is desirable for the Web server to include some geographic-specific information in this initial Web page. However, in many cases, there will be no conventional way to identify where the requesting computer is located geographically. For example, this may be the first time that the requesting computer has visited the Web site (meaning that no cookies previously could have been saved) or the requesting computer may have disabled the storage of cookies.

As indicated above, the present invention addresses this problem by estimating the number of hops made by the received data packet and transmitting one or more

probe packets based on this information. For example, if the Web server in the present example estimates that the incoming packet made 10 hops, then the probe packet might be designed to elicit information from the router that the probe packet encounters after 9 hops. Alternatively, several probe packets might be transmitted, designed to elicit information from the routers located 7-9 hops from the Web server (which presumably would be in the vicinity of the requesting computer).

By identifying the router or routers that are geographically closest to the requesting computer, the present invention can often provide at least a rough approximation as to the geographic location of such computer. Moreover, by designing the probe packets to investigate only the routers closest to the source of the received data packet, the number of probe packets required often can be minimized. As a result, the desired geographic information frequently can be obtained very quickly. In many cases, this speed can enable a Web server to transmit geographic-specific information in the initial Web page that is downloaded to a requesting computer. As noted above, absent special circumstances, this generally is not possible using conventional techniques.

Thus, independent claims 1, 31 and 33 are directed to methods, apparatuses and techniques for use by a first node on a network in determining the geographic location of a second node on the network. Initially, a data packet is received over the network from the second node. The data packet includes a network identifier for the second node and a Time-To-Live (TTL) field having a value that indicates a maximum additional number of hops that could have been made by the data packet. In response to this data packet, a probe packet, addressed to the network identifier for the second

node, is sent. The probe packet includes a TTL field having an initial value that is set based upon the value for the TTL field of the received data packet.

The foregoing combination of features is not disclosed or suggested by the applied art. In particular, the applied art does not disclose or suggest at least the feature of sending a probe packet addressed to the network identifier for a node from which a data packet was received, with the probe packet including a TTL field whose initial value is set based on the value for the TTL field of the received data packet.

In this regard, the following description summarizes Perkins's technique, as best understood by Applicant. Basically, Perkins concerns a technique for obtaining path redundancy over a packet network such as the Internet. Perkins indicates that his disclosed techniques are particularly applicable to voice-over-IP, where it is particularly important that packet loss not occur. See, e.g., column 1 lines 29-45. The assumption underlying Perkins's approach is that packet loss over the Internet tends to be path-related. See, e.g., column 1 lines 29-33. At the same time, however, the path that a stream of data packets will take generally cannot be known in advance. Therefore, Perkins addresses the problem by attempting to set up two simultaneous paths over the Internet, thereby providing a certain amount of redundancy. An example is shown in the Hops Digraph of Perkins's Figure 21.

To implement this technique, Perkins suggests the use of a plurality of proxy servers distributed throughout the Internet. Then, when an appropriately configured source node on the Internet attempts to set up a connection with a desired destination node, it first queries a "list server" to obtain destination addresses for two of the proxy servers. Based upon the network locations of the source node and the destination

node, the list server attempts to identify two proxy servers that would be on separate paths between the source node and the destination node. Once the source node obtains the network addresses for these two proxy servers, it directs its communications to the proxy servers (i.e., the destination field of each data packet is addressed to one of the proxy servers), rather than to the ultimate destination node. The body of each such data packet, however, includes the address for the destination node. Accordingly, when the proxy server receives the data packet it can automatically forward it to the destination node.

In this manner, the packet stream is routed through each proxy server. By simultaneously transmitting the same data stream through each proxy server, two redundant communication paths are established. As a result, even if packet loss occurs in one of the paths, the subject data usually will reach the destination node through the other path.

With regard to the present rejection, the Office Action asserts that column 31 lines 31-53 of Perkins discloses the receipt of a data packet over a network from a second node. However, it is unclear what the Office Action is asserting as the "first node", what it is asserting as the "second node", and what it is asserting as the "data packet" in this portion of Perkins. Rather, this portion of Perkins is only understood to discuss the Hops Digraph, the Link Matrix and the Hops Table of Figure 21; all of which provide information about the connectivity of the network. As discussed in this portion of Perkins, such tools can be used by the list server to help identify proxy servers for communications between a given source node and a given destination node. However, there appears to be no discussion of actually receiving a data packet and, therefore, no

indication as to which of Perkins's elements the present claim limitations supposedly correspond.

Similarly, the Office Action asserts that column 31 lines 27-30 of Perkins discloses that the "received data packet" includes both a network identifier and a TTL field. However, this portion of Perkins has been reviewed in detail and is not seen to mention any such data packet. Rather, this portion of Perkins merely discusses the Link Matrix which, as noted above, merely characterizes interconnections on the network (e.g., the Internet).

Lastly, the Office Action asserts that column 32 lines 47-54 of Perkins discloses the feature of sending a probe packet addressed to the network identifier for the "second node". However, this portion of Perkins is only seen to generally discuss one conceptual technique for use by the list server in identifying proxy servers. Specifically, the concept here is to: divide a region (or the entire world) into different geographic sections, somehow determine the sections in which the source and destination node lie, and then look up the pair of proxy servers based on this information. Nothing in this portion of Perkins appears to even remotely suggest sending a probe packet in any manner whatsoever, much less to a "second node" from which a data packet was received.

The Office Action acknowledges that Perkins does not say anything about setting an initial value for the TTL field of the probe packet based on the value of the TTL field of the received data packet. In order to make up for this deficiency, the Office Action cites column 3 lines 8-15 of Katsume. However, Katsume generally, and the cited portion of Katsume in particular, concerns techniques that are significantly different than

what is recited in the present claims, and also are significantly different than Perkins's technique.

More specifically, the cited portion of Katsume is taken from Katsume's Background section. A careful reading of its entire Background section indicates that Katsume concerns hop-count management in a label-switched network path. As discussed therein, a label-switched path is a defined path (e.g., over the Internet) through which data packets can be routed more efficiently than by utilizing conventional packet track transfer operations based on the packet's header information. See, e.g., column 1 lines 21-26. Katsume notes that each label-switched network path generally has one ingress node (where the packet enters the path), one egress node (where the packet exits the path), and one or more intermediate nodes (that transfer the packet based on input and output labels).

The paragraph preceding the cited portion of Katsume notes that one technique for managing hop count over a label-switched network path is for the ingress node to simply decrement the TTL value by the stored hop-count value for the entire label-switched path. The portion of Katsume that is cited then notes that such a technique is effective only when all of the intermediate nodes are disabled from updating the TTL. Clearly, nothing in this cited portion of Katsume pertains to setting the initial value for the TTL field of a probe packet based on the value for the TTL field of a received data packet.

In short, both Perkins and Katsume are significantly different than the presently recited claims and, in fact, are significantly different from each other. Many of the features of the present invention are missing from these references, regardless of

whether such references are considered individually or in combination. Still further, because they are significantly different from each other, it appears that there would have been no motivation to combine Perkins and Katsume in any manner whatsoever.

Based on the foregoing differences between the present claims and the applied art, independent claims 1, 31 and 33 are believed to be clearly allowable over the applied art.

Independent claims 18, 32 and 34 are directed to methods, apparatuses and techniques for use by a first node on a network in determining the geographic location of a second node on the network. Initially, a data packet is received from the second node, such data packet having arrived at the first node via an inbound path defined by an ordered sequence of routers. The number of hops made by the data packet is estimated based on information contained within the data packet. Then, probe packets are transmitted, such probe packets having been designed, based upon such estimated number of hops to elicit responses from a group of network devices that primarily includes the first few routers on the inbound path.

Based on the discussion above, it is apparent that no combination of Perkins and Katsume would have disclosed or suggested at least the foregoing features of: estimating the number of hops made by a received data packet based on information contained within the data packet and then transmitting probe packets having been designed, based upon such estimated number of hops, to elicit responses from a group of network devices that primarily includes the first few routers on the inbound path.

Accordingly, independent claims 18, 32 and 34 also are believed to be allowable over the applied art. The other claims in the application depend from the independent

claims discussed above, and are therefore believed to be allowable for at least the same reasons. In addition, each such dependent claim recites an additional feature of the invention that further distinguishes the invention from the applied art. Accordingly, the individual reconsideration of each on its own merits, particularly in view of the above remarks, is respectfully requested.

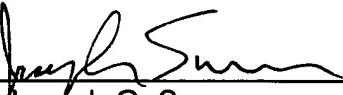
In view of the foregoing remarks, the entire application is believed to be in condition for allowance, and an indication to that effect is respectfully requested.

If there are any fees due in connection with the filing of this paper that have not been accounted for in this paper or the accompanying papers, please charge the fees to our Deposit Account No. 13-3735. If an extension of time under 37 C.F.R. 1.136 is required for the filing of this paper and is not accounted for in this paper or the accompanying papers, such an extension is requested and the fee (or any underpayment thereof) should also be charged to our Deposit Account No. 13-3735. A duplicate copy of this page is enclosed for that purpose.

Respectfully submitted,

MITCHELL, SILBERBERG & KNUPP LLP

Dated: May 4, 2004

By _____

Joseph G. Swan
Registration No. 41,338

MITCHELL, SILBERBERG & KNUPP LLP
11377 West Olympic Boulevard
Los Angeles, California 90064
Telephone: (310) 312-2000
Facsimile: (310) 312-3100